arXiv:1205.0679v1 [cs.LO] 3 May 2012

Lower Bounds for Existential Pebble Games and k-Consistency Tests

Christoph Berkholz

March 3, 2013

The existential k-pebble game characterizes the expressive power of the existential-positive k-variable fragment of first-order logic on finite structures. The winner of the existential k-pebble game on two given finite structures can easily be determined in time $O(n^{2k})$. We show that there is no $O(n^{(k-3)/12})$ time algorithm that decides which player can win the existential k-pebble game on two given structures. This lower bound is unconditional and does not rely on any unproven complexity theoretic assumptions.

Establishing strong k-consistency is a well-known heuristic for solving the constraint satisfaction problem (CSP). By the game characterization of Kolaitis and Vardi [13] our result implies that there is no $O(n^{(k-3)/12})$ time algorithm that decides if strong k-consistency can be established for a given CSP-instance.

1 Introduction

For two finite relational structures **A** and **B** the homomorphism problem asks if there is a mapping from the domain A of **A** to the domain B of **B** that preserves all relations. As pointed out by Feder and Vardi [6] this problem is equivalent to the constraint satisfaction problem (CSP) where the variables correspond to the domain of **A**, the values correspond to the domain of **B** and the constraints are encoded in the relations of **A** and **B**. Thus, every homomorphism from **A** to **B** corresponds to a solution of the CSP. Since the homomorphism problem and the CSP are NP-complete in general, there is need to look for heuristics. One well-known method introduced in the context of constraint satisfaction is the procedure of establishing strong k-consistency, which can be implemented by an $O(n^{2k})$ time algorithm (see e.g. [2,3,13]). We will explain this concept below in the setting of the homomorphism problem.

From a logical point of view, there is a homomorphism between two finite structures **A** and **B** if and only if every existential-positive first-order sentence that is true on **A** is also true on **B**. Instead of considering the full existential-positive fragment of first-order logic one can relax that question and ask whether every k-variable existential-positive first-order sentence true on **A** is also true on **B**. Such questions can be analyzed using combinatorial games. For the k-variable existential-positive fragment of first-order logic the corresponding game is the existential k-pebble game [11], which is defined in Section 2.

Kolaitis and Vardi [13] showed that this logical relaxation of the homomorphism problem is equivalent to the k-consistency heuristic. That is, strong k-consistency can be established if and only if Duplicator wins the existential k-pebble game on **A** and **B**. To sum up, the following three statements are equivalent on finite relational structures **A** and **B**:

- Strong k-consistency can be established.
- For every existential-positive k-variable first-order sentence φ: $\mathbf{A} \models \varphi \Longrightarrow \mathbf{B} \models \varphi$.

- Duplicator has a winning strategy in the existential k-pebble game.

Now we state our main result that provides a lower bound on the computational complexity of the statements above.

Theorem 1. *Given two finite relational structures* $\boldsymbol{A}$ *and* $\boldsymbol{B}$ *and an integer* $k \geq 15$*, deciding if Duplicator has a winning strategy on* $\boldsymbol{A}$ *and* $\boldsymbol{B}$ *in the existential* k*-pebble game requires time* $\omega((\|\boldsymbol{A}\| + \|\boldsymbol{B}\|)^{\frac{k-3}{12}})$.

From this theorem we directly get an $\omega(n^{\frac{k-3}{12}})$ lower bound for deciding if strong k-consistency can be established, where n is the size of the CSP-instance. As an upper bound, the query "Does Spoiler win the existential k-pebble game on $\mathbf{A}$ and $\mathbf{B}$?" is LFP^{2k}-definable for two given structures $\mathbf{A}$ and $\mathbf{B}$ [12] and is decidable by an $O(|A|^k|B|^k)$ time algorithm.

We prove Theorem 1 by a reduction from the k-pebble game of Kasai, Adachi and Iwata [9], called KAI-game, to the existential $(k+1)$-pebble game. Our result then follows from the $\omega(n^{\frac{k-2}{4}})$ lower bound for this game [1], which in turn follows from the deterministic time hierarchy theorem.

1.1 Related Work

Kolaitis and Panttaja [10] proved that for every fixed $k \geq 2$ the problem of determining the winner of the existential k-pebble game is complete for PTIME under LOGSPACE reductions. Furthermore, they established that the problem is complete for EXPTIME when k is part of the input. It follows that there is no algorithm for this problem whose running time is polynomial in the size of the structures as well as in the number of pebbles. Parameterized by the number of pebbles k, the problem is known to be W[1]-hard. This follows directly from the fact that a graph G contains a k-clique if and only if Duplicator has a winning strategy for the existential k-pebble game on the complete graph on k vertices and G. Thus, the existence of an algorithm of running time $f(k)n^c$ for some computable function f and constant c would imply W[1] = FPT, an unlikely event in parameterized complexity theory.

However, since we do not know whether W[1] = FPT it is consistent with our previous knowledge that there exists an $O(2^k n^2)$ algorithm determining the winner of the existential k-pebble game on two relational structures. Thus, for every fixed k, it was possible that there exists a quadratic time algorithm deciding if strong k-consistency can be established.

To prove the EXPTIME-completeness Kolaitis and Panttaja reduced the KAI-game to the existential pebble game. In this reduction the number of pebbles used in the existential pebble game depends on the size of the KAI-game instance and is not bounded by any function on the number of pebbles used in the KAI-game. Thus, their reduction fails to prove a lower bound for fixed k and it was left as an open question if such a lower bound can be proven. In this paper we reduce the k-pebble KAI-game to the existential $(k+1)$-pebble game, and thus, keep the parameter small. Some constructions are quite similar to those used by Kolaitis and Panttaja. However, the proof differs significantly at crucial points. Furthermore, to devise the winning strategy for Duplicator we use a different proof technique.

Finally, it should be noticed that there has recently been interest in the parameterized complexity of k-consistency. Gaspers and Szeider [7] showed that the problem "given two finite relational structures, are they strongly k-consistent?" is complete for co-W[2] when k is the parameter. By contrast, we prove that the problem "given two finite relational structures, can they be made strongly k-consistent?" is complete for XP when k is the parameter.

1.2 Further Applications in Finite Model Theory

We can improve Theorem 1 in two ways. First, all partial homomorphisms in Duplicator's winning strategy (defined below) are in fact partial isomorphisms. Thus, Duplicator has a winning strategy in the k-pebble game that corresponds to the existential k-variable fragment of first-order logic, where negation is allowed in front of atomic formulas. This implies that it requires $\omega((\|\mathbf{A}\| + \|\mathbf{B}\|)^{\frac{k-3}{12}})$ time to decide if every existential k-variable first-order sentence true on **A** is also true on **B**. Second, the structures our reduction constructs are directed graphs. Therefore, Theorem 1 holds even when we restrict ourselves to σ-structures, where σ is a relational signature containing at least one binary relation.

1.3 Establishing Strong k-Consistency

To fix the terminology, we briefly introduce the concept of establishing strong k-consistency as it is defined in [13]. Let **A** and **B** be two finite relational structures with universes A and B. A *k-partial homomorphism* is a partial homomorphism with domain size k. **A** and **B** are *k-consistent* if for every (k-1)-partial homomorphism h from **A** to **B** and every $a \in A$ there is a partial homomorphism that extends h and is defined on a. Two structures are *strongly k-consistent* if they are i-consistent for every $i \leq k$.

Strong k-consistency *can be established* for **A** and **B** if there are two strongly k-consistent structures **A'** and **B'** over the same universes A and B such that the following two statements hold.

- Every k-partial homomorphism from **A'** to **B'** is a k-partial homomorphism from **A** to **B**.
- Every function from A to B is a homomorphism from **A'** to **B'** if and only if it is a homomorphism from **A** to **B**.

Loosely speaking, strong k-consistency can be established for two structures if they can be made strongly k-consistent by adding new relations and without changing the solution space with respect to the homomorphism problem. It is easy to see that if there is a homomorphism from **A** to **B**, then strong k-consistency can be established. Although the converse is not true in general, it holds for some classes of structures [2, 4].

2 Definitions and Proof of Theorem 1

An instance of the k-pebble KAI-game (introduced in [9] as pebble game of fixed rank) is a tuple $(X, R, \mathfrak{s}, \gamma)$, where X is the universe, $R = R' \times \binom{[k]}{2} \subseteq X^3 \times [k]^2$ the set of rules, $\mathfrak{s}\colon [k] \to X$ the start position and $\gamma \in X$ the goal. We let $[k]$ be the set of k pebbles in the game. A rule is of the form (u, v, w, c, d), with $c \neq d$, $u \neq v \neq w$ and has the intended meaning that if pebble c is on u and pebble d is on v and there is no pebble on w then one player can move pebble c from u to w. This is a slightly more wasteful notion as the original one used in [9], where the relation $R' \subseteq X^3$ (instead of $R' \times \binom{[k]}{2}$) is given as input. However, this technical modification does not affect the purpose of the game and increases the size of an instance only by a constant factor, if k is fixed, and by a polynomial factor, if k is part of the input. A *position* of the KAI-game is an injective mapping $\mathfrak{p}\colon [k] \to X$. A rule $r = (u, v, w, c, d) \in R$ is *applicable* to a position $\mathfrak{p}$ if $\mathfrak{p}(c) = u$, $\mathfrak{p}(d) = v$ and $\mathfrak{p}(z) \neq w$ for all $z \in [k]$. Furthermore, if r is applicable to $\mathfrak{p}$ then $r(\mathfrak{p})$ denotes the position defined as $r(\mathfrak{p})(c) = w$ and $r(\mathfrak{p})(z) = \mathfrak{p}(z)$, for all $z \in [k] \setminus \{c\}$. The set of all rules in R applicable to a position $\mathfrak{p}$ is denoted by $\mathrm{appl}(\mathfrak{p})$, and $T_r(\mathfrak{p}) \subseteq [k]$ is the set of KAI-pebbles i such that $\mathfrak{p}(i)$ contradicts the applicability condition of rule r.

The k-pebble KAI-game is played by two players and proceeds in rounds. In the first round Player 1 starts with position $\mathfrak{s}$ and chooses a rule $r \in \mathrm{appl}(\mathfrak{s})$. The new position is $\mathfrak{p} = r(\mathfrak{s})$. In the next round Player 2 chooses a rule $r \in \mathrm{appl}(\mathfrak{p})$ and applies it to $\mathfrak{p}$. Then it is Player 1's turn and so on. Player 1 wins the game if he reaches a position $\mathfrak{p}$, where $\mathfrak{p}(z) = \gamma$ for one $z \in [k]$ or where Player 2 is unable to move. Player 2 wins if she has a strategy ensuring that Player 1 cannot reach such a position. The following definition makes that formal.

Definition 2. A winning strategy for Player 2 in the k-pebble KAI-game is a set $\mathcal{K} \subseteq \{(\mathfrak{p}, i) \mid \mathfrak{p}\colon [k] \to X, i \in \{1, 2\}\}$ such that

- $(\mathfrak{s}, 1) \in \mathcal{K}$,
- for all $(\mathfrak{p}, 1) \in \mathcal{K}$ and all $r \in \operatorname{appl}(\mathfrak{p})$: $(r(\mathfrak{p}), 2) \in \mathcal{K}$,
- for all $(\mathfrak{p}, 2) \in \mathcal{K}$ there exists an $r \in \operatorname{appl}(\mathfrak{p})$ such that $(r(\mathfrak{p}), 1) \in \mathcal{K}$,
- for all $(\mathfrak{p}, 2) \in \mathcal{K}$ and all $j \in [k]$: $\mathfrak{p}(j) \neq \gamma$.

Kasai, Adachi and Iwata [9] showed that the problem of determining the winner of the k-pebble KAI-game is PTIME-complete (for every fixed $k \geq 3$) under LOGSPACE-reductions and complete for EXPTIME when k is part of the input. Furthermore, they showed that the problem to requires $\Omega(N^{\frac{k-1}{4}-\varepsilon})$ time on multi-tape Turing machines for every fixed $k \geq 6$ and any $\varepsilon > 0$, where N is the size of the input [1]. This proof also leads to XP-completeness of the k-pebble KAI-game with parameter k [5].

The existential (k+1)-pebble game [11] is played by two players Spoiler and Duplicator on two relational structures **A** and **B** with domains A and B, respectively. First, Spoiler puts pebbles $a_1, \dots, a_{k+1}$ on elements of A and Duplicator answers by putting pebbles $b_1, \dots, b_{k+1}$ on elements of B. In each further round Spoiler picks up a pebble a_i from A and places it on another element in A and Duplicator moves the corresponding pebble b_i in B. Spoiler wins the game if he can reach a position where the mapping defined by $a_i \mapsto b_i$ is not a partial homomorphism from **A** to **B**. Duplicator wins the game if she has a winning strategy that tells her for every move of Spoiler how to place her pebbles such that the positions of the pebbles define a partial homomorphism. A winning strategy for Duplicator can be stated formally as a set of partial homomorphisms:

Definition 3. A *winning strategy* for Duplicator in the existential (k+1)-pebble game on structures **A** and **B** is a nonempty family $\mathcal{H}$ of partial homomorphisms from **A** to **B** satisfying the following properties:

- (closure) If $h \in \mathcal{H}$ and $g \subset h$ then $g \in \mathcal{H}$.
- (extension) For every $g \in \mathcal{H}$, $|\operatorname{Dom}(g)| \leq k$, and every $z \in A$ there is an $h \in \mathcal{H}$ with $g \subseteq h$ and $z \in \operatorname{Dom}(h)$.

For a set H of partial homomorphisms from **A** to **B** we let $\operatorname{cl}(H) := \{g \mid g \subseteq h, h \in H, |\operatorname{Dom}(g)| \leq k+1\}$ and write $\operatorname{cl}(h)$ instead of $\operatorname{cl}(\{h\})$. It is easy to see that if h is a total homomorphism from **A** to **B**, then $\operatorname{cl}(h)$ is a winning strategy in the existential (k+1)-pebble game on **A** and **B**. Now we state our main lemma and prove Theorem 1.

Lemma 4 (Main Lemma). *There is a reduction from the k-pebble KAI-game to the existential (k+1)-pebble game that computes for every instance* $I = (X, R, \mathfrak{s}, \gamma)$ *two directed graphs* G_S *and* G_D *such that the following constraints hold:*

- *Player 1 has a winning strategy in the k-pebble KAI-game on I if and only if Spoiler has a winning strategy in the existential (k+1)-pebble game on* G_S *and* G_D.
- $|V(G_S)| + |V(G_D)| = O(|X| \cdot |R| \cdot k^2)$.
- $|E(G_S)| + |E(G_D)| = O(k^4(|X|^2|R| + n|R|^2))$.
- *The reduction is computable in* $\mathrm{DTIME}(O(\|I\|^3))$ *and in* LOGSPACE.

Proof of Theorem 1. Let $k \geq 15$ be a fixed integer. Assume that $\mathbb{A}$ is an algorithm that determines the winner of the existential k-pebble game on structures $\mathbf{A}$ and $\mathbf{B}$ in time $O((\|\mathbf{A}\|+\|\mathbf{B}\|)^{\frac{k-3}{12}})$. Let $\mathbb{B}$ be the algorithm that first applies the reduction from Lemma 4 to a given instance I of the (k-1)-pebble KAI-game and then executes $\mathbb{A}$. Since $\|G_S\|+\|G_D\| = O(\|I\|^3)$, $\mathbb{B}$ has running time $O(\|I\|^3+\|I\|^{\frac{k-3}{4}})$, and thus solves the k'-pebble KAI-game in time $O(\|I\|^{\frac{k'-2}{4}})$. This contradicts the $\Omega(N^{\frac{k'-1}{4}-\varepsilon})$ lower bound [1]. □

In addition, Lemma 4 also implies EXPTIME-completeness when k is part of the input, PTIME-completeness for every fixed $k \geq 4$ and XP-completeness when k is the parameter.

In our reduction we first construct two colored graphs out of smaller graphs, called gadgets. In order to prove the existence of a winning strategy for one player, we combine strategies for the gadgets to a strategy for the whole graph. The easier part is to do that for Spoiler. As in [8] and [10], we say that Spoiler *can reach* position p_j from position p_i if he has a strategy in the game such that starting from position p_i he either wins the game, or position p_j occurs in the game after some finite number of rounds. Since this relation is transitive, we can combine such strategies to show that Spoiler can reach some position p from $\emptyset$; if p does not define a partial homomorphism, this will give us a winning strategy for Spoiler.

For Duplicator this is more difficult. A *critical strategy* in the existential (k+1)-pebble game is a nonempty family $\mathcal{H}$ of partial homomorphisms satisfying the closure property (Definition 3) together with a set of *critical positions* $\operatorname{crit}(\mathcal{H}) \subset \mathcal{H}$ such that $h \in \operatorname{crit}(\mathcal{H}) \implies |\operatorname{Dom}(h)| = k$ and all $g \in \mathcal{H} \setminus \operatorname{crit}(\mathcal{H})$ satisfy the extension property. A critical strategy is nearly a winning strategy in the sense that Duplicator wins unless the game reaches a critical position. Note that a critical strategy with $\operatorname{crit}(\mathcal{H}) = \emptyset$ is a winning strategy and every critical strategy in the (k+1)-pebble game is a winning strategy in the k-pebble game. Let $\hat{\mathcal{H}} := \mathcal{H} \setminus \operatorname{crit}(\mathcal{H})$. As for winning strategies, the union of critical strategies is also a critical strategy. The following lemma enables us to construct a winning strategy out of critical strategies.

Lemma 5. *If $\mathcal{H}_1, \ldots, \mathcal{H}_l$ are critical strategies on the same structures and for all $i \in [l]$ and all $p \in \operatorname{crit}(\mathcal{H}_i)$ there exists a $j \in [l]$ such that $p \in \hat{\mathcal{H}}_j$, then $\bigcup_{i\in[l]} \mathcal{H}_i$ is a winning strategy on these structures.* □

Every gadget Q that we construct consists of two graphs Q_S and Q_D for Spoiler's and Duplicator's side, respectively. The graphs contain *boundary vertices* $\operatorname{bd}(Q_S) \subseteq V(Q_S)$ and $\operatorname{bd}(Q_D) \subseteq V(Q_D)$, which are the vertices where the gadget is connected to other gadgets. That is, vertices in $V(Q_S) \setminus \operatorname{bd}(Q_S)$ ($V(Q_D) \setminus \operatorname{bd}(Q_D)$) are only adjacent to vertices in $V(Q_S)$ ($V(Q_D)$). A *boundary of a strategy* $\mathcal{H}$ on a gadget Q is a mapping $\beta\colon \operatorname{bd}(Q_S) \to \operatorname{bd}(Q_D)$ such that $h(z) = \beta(z)$ for all $h \in \mathcal{H}$ and all $z \in \operatorname{bd}(Q_S) \cap \operatorname{Dom}(h)$. We say that two strategies $\mathcal{G}$ and $\mathcal{H}$ on gadgets $Q^{\mathcal{G}}$ and $Q^{\mathcal{H}}$ are *connectable*, if they have boundaries $\beta_{\mathcal{G}}$ and $\beta_{\mathcal{H}}$ and it holds that $\beta_{\mathcal{G}}(z) = \beta_{\mathcal{H}}(z)$ for all $z \in \operatorname{bd}(Q_S^{\mathcal{G}}) \cap \operatorname{bd}(Q_S^{\mathcal{H}})$. If $\mathcal{G}$ and $\mathcal{H}$ are two connectable strategies, we define

$$\mathcal{G} \uplus \mathcal{H} = \{g \cup h \mid g \in \mathcal{G}, h \in \mathcal{H}\}.$$

Lemma 6. *Let $\mathcal{G}$ and $\mathcal{H}$ be two connectable critical strategies on gadgets $Q^{\mathcal{G}} = (Q_S^{\mathcal{G}}, Q_D^{\mathcal{G}})$ and $Q^{\mathcal{H}} = (Q_S^{\mathcal{H}}, Q_D^{\mathcal{H}})$. The composition $\mathcal{G} \uplus \mathcal{H}$ is a critical strategy on $Q_S^{\mathcal{G}} \cup Q_S^{\mathcal{H}}$ and $Q_D^{\mathcal{G}} \cup Q_D^{\mathcal{H}}$ with* $\operatorname{crit}(\mathcal{G} \uplus \mathcal{H}) = \operatorname{crit}(\mathcal{G}) \cup \operatorname{crit}(\mathcal{H})$. □

We use the operator $\uplus$ to construct global critical strategies for the whole graph out of critical strategies on the gadgets. Then we show that the union of those global critical strategies is by Lemma 5 a winning strategy for Duplicator.

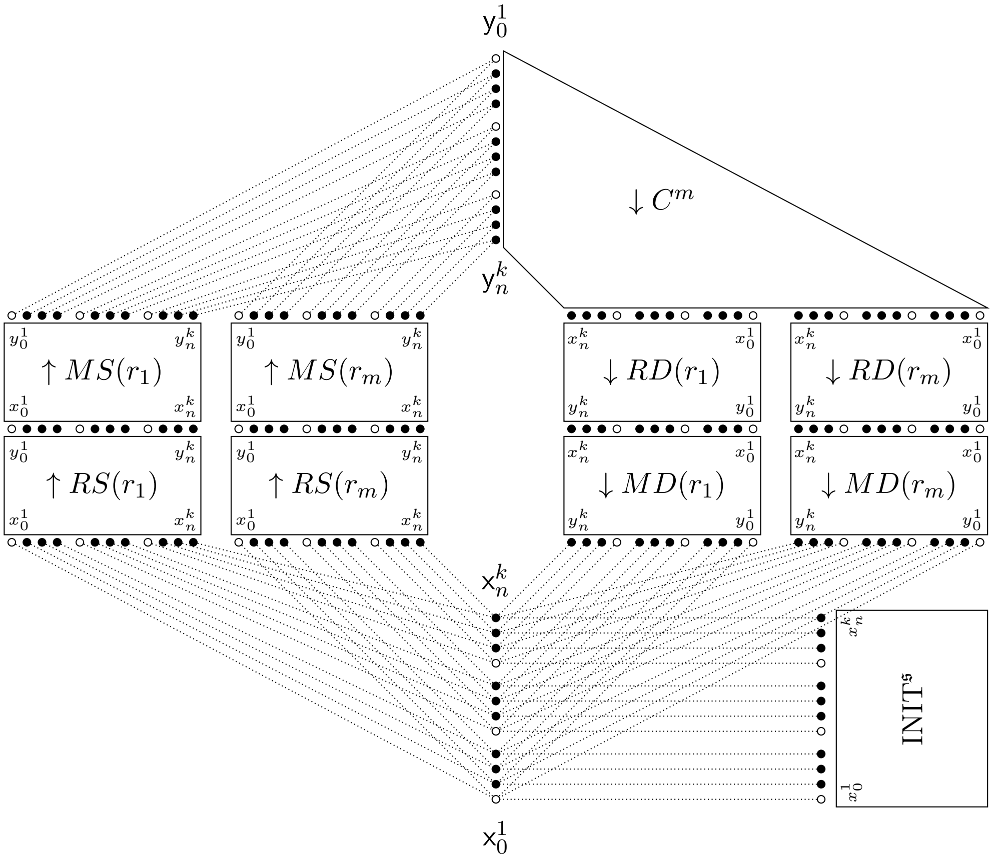

Figure 1: The graph G_D. The dotted edges are only for visual purposes and need to be contracted.

3 The Construction

Let $([n], R, \mathfrak{s}, \gamma)$ be an instance of the k-pebble KAI-pebble game and $m := |R|$. As in [10], the main idea is to simulate every run of the KAI-game within the existential pebble game such that Spoiler imitates the moves of Player 1 and Duplicator imitates the moves of Player 2. First, we construct two colored simple graphs, G_S and G_D, and then show how to omit the colors while switching to directed graphs. We color every vertex of Spoiler's graph G_S with a different color and partition the vertices of Duplicator's graph with these colors. Thus, whenever Spoiler pebbles a vertex in G_S there is a corresponding set of vertices in G_D Duplicator can pebble.

To encode a position of the KAI-game in the existential $(k{+}1)$-pebble game we introduce the vertices $\{\mathsf{x}^1, \ldots, \mathsf{x}^k\}$ in Spoiler's graph and $\{\mathsf{x}^i_l \mid i \in [k], 0 \leq l \leq n\}$ in Duplicator's graph. For each i, all the vertices $\{\mathsf{x}^i\} \cup \{\mathsf{x}^i_0, \ldots, \mathsf{x}^i_n\}$ are colored with the same unique color, denoted by c_{x^i}. The vertices x^i_0 play a special role in the construction, so we will draw $\circ$ for vertices with subindex 0 in the figures and $\bullet$ for all other vertices. Furthermore, we introduce vertices y^i,y^i_l in an appropriate manner.

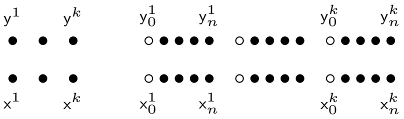

If $\mathfrak{p}\colon [k] \to [n]$ is a position of the k pebbles in the KAI-game and it is Player 1's turn, then $\{(\mathsf{x}^i, \mathsf{x}^i_{\mathfrak{p}(i)}) \mid i \in [k]\}$ is the corresponding position in the existential pebble game. If it is Player 2's turn, then $\{(\mathsf{y}^i, \mathsf{y}^i_{\mathfrak{p}(i)}) \mid i \in [k]\}$ is the corresponding position. During the course of the game Spoiler pebbles some vertex x^i asking, "Where does KAI-pebble i lie on?" Due to the coloring Duplicator has to answer with some vertex x^i_l meaning "KAI-pebble i lies on node l." The vertices $\circ$ are used to handle the case when Spoiler does not play in the intended way, that is, Spoiler has a winning strategy if and only if he has a winning strategy on the $\bullet$ vertices. In order to name positions that include $\circ$ vertices, we define for positions $\mathfrak{p}$ and sets $T \subseteq [k]$ the mapping $(\mathfrak{p}, T)$ as

$$(\mathfrak{p}, T)(i) = \begin{cases} \mathfrak{p}(i),\ i \notin T, \\ 0, \text{ else,} \end{cases}$$

and write $\mathfrak{p}$ for $(\mathfrak{p}, \emptyset)$ and $\mathbf{0}$ for $(\mathfrak{p}, [k])$. Now we have to introduce various gadgets to ensure the following:

- For the start position $\mathfrak{s}$, Spoiler can reach $\{(\mathsf{x}^i, \mathsf{x}^i_{\mathfrak{s}(i)}) \mid i \in [k]\}$ from $\emptyset$.
- For a position $\mathfrak{p}$ and every rule $r \in \mathrm{appl}(\mathfrak{p})$, Spoiler can reach $\{(\mathsf{y}^i, \mathsf{y}^i_{r(\mathfrak{p})(i)}) \mid i \in [k]\}$ from $\{(\mathsf{x}^i, \mathsf{x}^i_{\mathfrak{p}(i)}) \mid i \in [k]\}$.
- Spoiler can reach $\{(\mathsf{x}^i, \mathsf{x}^i_{r(\mathfrak{p})(i)}) \mid i \in [k]\}$ from $\{(\mathsf{y}^i, \mathsf{y}^i_{\mathfrak{p}(i)}) \mid i \in [k]\}$ for a rule $r \in \mathrm{appl}(\mathfrak{p})$ of Duplicator's choice.
- $\{(\mathsf{y}^i, \mathsf{y}^i_{\mathfrak{p}(i)}) \mid i \in [k]\}$ is not a partial homomorphism if $\mathfrak{p}(i) = \gamma$ for some $i \in [k]$.

We give a brief description of the construction. Duplicator's graph is illustrated in Figure 1. The gagdets are glued together at their boundary vertices and the vertex blocks that are glued together inherit their colors. In order to make sure that the colors partition the graphs we define a new color for every combination of colors occurring at one vertex in the graph. In Spoiler's graph we do this analog with Spoiler's side of the gadgets.

To implement the last condition, we simply delete the color c_{y^i} from y^i_γ for all $i \in [k]$. Since y^i and all y^i_l, $l \in [n] \setminus \{\gamma\}$, are still colored c_{y^i}, it follows that the mapping $\mathsf{y}^i \mapsto \mathsf{y}^i_l$ that encodes

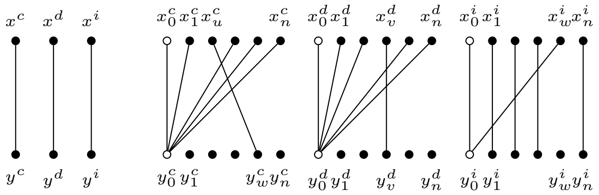

Figure 2: Rule Gadget $RS(u, v, w, c, d)$ The range of i is $[k] \setminus \{c, d\}$.

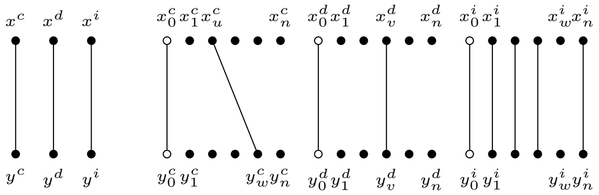

Figure 3: Rule Gadget $RD(u, v, w, c, d)$. The range of i is $[k] \setminus \{c, d\}$.

"KAI-pebble i lies on node l" is a partial homomorphism if and only if l is not the goal node γ. To make sure that Spoiler can reach the start position, we introduce the Initialization Gadget $\mathrm{INIT}^{\mathfrak{s}}$, whose boundary $x^1, \dots, x^k$ in Spoiler's graph and $x_0^1, \dots, x_n^k$ in Duplicator's graph is identified with vertices $\mathsf{x}^1, \dots, \mathsf{x}^k$ in Spoiler's graph and $\mathsf{x}_0^1, \dots, \mathsf{x}_n^k$ in Duplicator's graph.

The boundary of the other gadgets has a similar form and can be divided in input vertices x (with certain indices) and output vertices y that are colored in the same way as the x-vertices. The direction of the gadgets is indicated in Figure 1 by arrows. Thus, the players are intended to move clockwise in the graph.

For each rule r we define different rule gadgets $RS(r)$ and $RD(r)$ in which Spoiler can reach the positions $(y^i, y^i_{r(\mathfrak{p})(i)})$ from $(x^i, x^i_{\mathfrak{p}})$ if r is applicable to $\mathfrak{p}$. The choice gadget C^m enables Duplicator to choose one of the m rules she wants to apply. Such a choice gadget is not necessary for Spoiler, because he can force Duplicator to play on the chosen rule gadget by the coloring of the vertices. The most complex gadget is the Multiple Input One-way Switch $M^{k,n}$, which is a generalization of the Multiple Input One-way Switch defined in [10]. In our construction we put one copy of $M^{k,n}$ at the output vertices of every Rule Gadget. On the one hand, this gadget ensures that Spoiler needs to pick up all $k+1$ pebbles to make progress through the switch. This is necessary because position $(\mathsf{x}^i, \mathsf{x}^i_{r'(r(\mathfrak{p}))(i)})$ occurs after two rounds in the KAI-game and Spoiler would win if he leaves a pebble pair on $(\mathsf{x}^i, \mathsf{x}^i_{\mathfrak{p}(i)})$. On the other hand only valid positions (where no $\circ$ vertex is pebbled) can be pebbled through the switch. Thus, pebbling on $\circ$ vertices is a method for Duplicator to penalize Spoiler for playing incorrectly.

3.1 Rule Gadgets

The Rule Gadgets $RS(r)$ and $RD(r)$ consist of input vertices $x^1, \dots, x^k$ in Spoiler's graph and $x_0^1, \dots, x_n^k$ in Duplicator's graph, and output vertices $y^1, \dots, y^k$ and $y_0^1, \dots, y_k^n$. For each rule $r = (u, v, w, c, d)$ we connect the vertices in the gadgets $RS(r)$ and $RD(r)$ as shown in figures 2 and 3.

If r is applicable to $\mathfrak{p}$, then Spoiler can reach the position $\{(y^i, y^i_{r(\mathfrak{p})(i)}) \mid i \in [k]\}$ from $\{(x^i, x^i_{\mathfrak{p}(i)}) \mid i \in [k]\}$ in both gadgets. To do so, he picks up the $(k+1)$st pebble and puts it on y^1. Then he picks up the pebble from x^1 and puts it on y^2 and so on. This fact is stated in Lemma 7(i) and 8(i). Assume that r is not applicable to $\mathfrak{p}$ (then $T_r(\mathfrak{p}) \neq \emptyset$) and the current position is $\{(x^i, x^i_{\mathfrak{p}(i)}) \mid i \in [k]\}$. On $RS(r)$ Duplicator can pebble some $\circ$ vertex y_0^i when Spoiler asks for y^i ($i \in T_r(\mathfrak{p})$), and thus can avoid valid positions on the y-vertices. Therefore, Spoiler is penalized when he chooses a rule r not applicable to $\mathfrak{p}$ and plays on $RS(r)$. This strategy is stated in Lemma 7(ii) for $T = \emptyset$. If Duplicator chooses a rule r not applicable to $\mathfrak{p}$ and plays on $RD(r)$, then she will be penalized, because Spoiler wins immediately from position $\{(x^i, x^i_{\mathfrak{p}(i)}) \mid i \in [k]\}$ (Lemma 8(iii)) by pebbling on y^i for some $i \in T_r(\mathfrak{p})$. Furthermore, Lemma 7(ii) and 8(ii) state that for invalid positions on the x-vertices (i.e. $T \neq \emptyset$), Duplicator can avoid valid positions on the y-vertices. To construct Duplicators strategies formally, consider the total homomorphisms that are defined by the requirements on the boundary.

Lemma 7. *For every rule $r = (u, v, w, c, d)$ and position $\mathfrak{p}\colon [k] \to [n]$ the following holds in the existential (k+1)-pebble game on $RS(r)$:*

(i) If $r \in \operatorname{appl}(\mathfrak{p})$, then Spoiler can reach $\{(y^i, y^i_{r(\mathfrak{p})(i)}) \mid i \in [k]\}$ from $\{(x^i, x^i_{\mathfrak{p}(i)}) \mid i \in [k]\}$.

(ii) Duplicator has a winning strategy $\mathcal{R}_{(\mathfrak{p},T)}$ with boundary $\{(x^i, x^i_{(\mathfrak{p},T)(i)}) \mid i \in [k]\} \cup \{(y^i, y^i_{(r(\mathfrak{p}),T \cup T_r(\mathfrak{p}))(i)}) \mid i \in [k]\}$, for all $T \subseteq [k]$.

Lemma 8. *For every rule $r = (u, v, w, c, d)$ and position $\mathfrak{p}\colon [k] \to [n]$ the following holds in the existential (k+1)-pebble game on $RD(r)$:*

(i) If $r \in \operatorname{appl}(\mathfrak{p})$, then Spoiler can reach $\{(y^i, y^i_{r(\mathfrak{p})(i)}) \mid i \in [k]\}$ from $\{(x^i, x^i_{\mathfrak{p}(i)}) \mid i \in [k]\}$.

(ii) If $r \in \operatorname{appl}(\mathfrak{p})$, then Duplicator has a winning strategy $\mathcal{R}_{(\mathfrak{p},T)}$ with boundary $\{(x^i, x^i_{(\mathfrak{p},T)(i)}) \mid i \in [k]\} \cup \{(y^i, y^i_{(r(\mathfrak{p}),T)(i)}) \mid i \in [k]\}$, for all $T \subseteq [k]$.

(iii) If $r \notin \operatorname{appl}(\mathfrak{p})$, then Spoiler wins from $\{(x^i, x^i_{\mathfrak{p}(i)}) \mid i \in [k]\}$.

3.2 The Multiple Input One-Way Switch

As for the Rule Gadgets, the Multiple Input One-way Switch has input vertices $x^1, \dots, x^k$ in Spoiler's graph and $x^1_0, \dots, x^k_n$ in Duplicator's graph, and output vertices $y^1, \dots, y^k$ and $y^1_0, \dots, y^k_n$, respectively. We say that a position $(\mathfrak{p}, T)$ is on the input (output) to denote the positions $\{(x^i, x^i_{(\mathfrak{p},T)(i)}) \mid i \in [k]\}$ ($\{(y^i, y^i_{(\mathfrak{p},T)(i)}) \mid i \in [k]\}$).

For Spoiler the switch ensures that he can reach $\mathfrak{p}$ on the output from $\mathfrak{p}$ on the input. For Duplicator there are several strategies. Duplicator has a winning strategy called *output strategy*, where any position is on the output and $\mathbf{0}$ is on the input. This ensures that Spoiler cannot move backwards and reach $\mathfrak{p}$ on the input from $\mathfrak{p}$ on the output. Next, for every nonempty $T \subseteq [k]$ Duplicator has a winning strategy where $(\mathfrak{p}, T)$ is on the input and $\mathbf{0}$ is on the output. Thus, she has a strategy such that Spoiler can reach only the $\mathbf{0}$ position on the output from invalid positions on the input. These strategies are called *restart strategies*. We will see later that Spoiler has to restart the game, that is, he has to pick up all pebbles and start playing on the Initialization Gadget, if he reaches a position that is contained in a restart strategy. To ensure that Spoiler picks up all pebbles when reaching $\mathfrak{p}$ on the output from $\mathfrak{p}$ on the input, Duplicator has a critical *input strategy* with $\mathfrak{p}$ on the input and $\mathbf{0}$ on the output, whose critical positions are either contained in an output strategy (where $\mathfrak{p}$ is on the output) or in a restart strategy. Thus, if Spoiler wants to make progress on the switch, he has to pebble a critical position in that input strategy using k pebbles and picking up the $(k+1)$st pebble. If he is clever, then he pebbles a critical position that is contained in the output strategy.

In order to define the *Multiple Input One-Way Switch* $M^{k,n}$ we construct the two graphs $M^{k,n}_S$ for Spoiler's graph and $M^{k,n}_D$ for Duplicator's graph. Let

$$\begin{aligned} V(M^{k,n}_S) &= \{x^i, a^i, b^i, y^i \mid i \in [k]\}, \\ E(M^{k,n}_S) &= \big\{\{x^i, a^i\}, \{b^i, y^i\} \mid i \in [k]\big\} \\ &\qquad\qquad \cup \big\{\{a^i, a^j\}, \{b^i, b^j\} \mid i, j \in [k]; i \neq j\big\} \end{aligned}$$

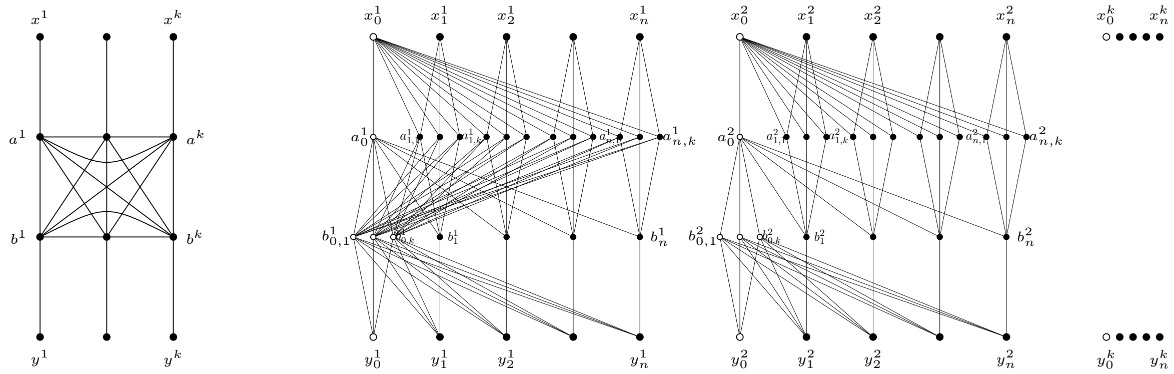

Figure 4: Subgraph of the Multiple Input One-way Switch $M^{k,n}$. For the first block in Duplicator's side, all inner-block edges are drawn.

$$\cup \big\{\{a^i, b^j\} \mid i,j \in [k]\big\}.$$

That is, $M_S^{k,n}$ simply consists of k input vertices x^i and k output-vertices y^i, which are each connected to one vertex of a $2k$-clique. For Duplicator's side of the graph, we define for $i \in [k]$:

$$\begin{aligned} X^i &= \{x_s^i \mid 0 \le s \le n\}, & Y^i &= \{y_s^i \mid 0 \le s \le n\} \\ A_+^i &= \{a_{s,l}^i \mid s \in [n], l \in [k]\}, & A^i &= A_+^i \cup \{a_0^i\} \\ B_+^i &= \{b_s^i \mid s \in [n]\}, & B^i &= B_+^i \cup \{b_{0,l}^i \mid l \in [k]\}. \end{aligned}$$

The set of vertices of $M_D^{k,n}$ is

$$V(M_D^{k,n}) = \bigcup_{i \in [k]} \left(X^i \cup A^i \cup B^i \cup Y^i\right).$$

The graph consists of k blocks, where the i-th block contains the vertices $X^i \cup A^i \cup B^i \cup Y^i$. For every $i \in [k]$ we color $\{x^i\} \cup X^i$ as well as $\{a^i\} \cup A^i$, $\{b^i\} \cup B^i$ and $\{y^i\} \cup Y^i$ with a unique color. We first define the inner-block edges E^i, which are also shown in Figure 4, and then the inter-block edges $E^{i,j}$ (for notational convenience we always assume $l, p \in [k]$ and $s, q \in [n]$):

$$\begin{aligned} E^i =& \big(\{x_0^i\} \times A^i\big) \cup \big\{\{x_s^i, a_{s,l}^i\}\big\} \cup \big(\{a_0^i\} \times B^i\big) \cup \\ & \big\{\{a_{s,l}^i, b_s^i\}\big\} \cup \big\{\{a_{s,l}^i, b_{0,p}^i\} \mid l \neq p\big\} \cup \big\{\{b_s^i, y_s^i\}\big\} \cup \\ & \big(\{b_{0,l}^i \mid l \in [k]\} \times Y^i\big), \end{aligned}$$

$$\begin{aligned} E^{i,j} =& \big\{\{a_{s,l}^i, a_{q,p}^j\} \mid l \neq p\big\} \cup \big\{\{a_0^i, a_{s,l}^j\}\big\} \cup \big\{\{b_{0,l}^i, b_{0,p}^j\}\big\} \cup \\ & \big\{\{b_s^i, b_q^j\}\big\} \cup \big\{\{a_0^i, b_{0,l}^j\}\big\} \cup \big\{\{a_0^i, b_s^j\}\big\} \cup \\ & \big\{\{a_{s,l}^i, b_q^j\}\big\} \cup \big\{\{a_{s,l}^i, b_{0,p}^j\} \mid l \neq p\big\}. \end{aligned}$$

Finally, $E(M_D^{k,n}) = \bigcup_{i \in [k]} E^i \cup \bigcup_{i,j \in [k]; i \neq j} E^{i,j}$.

Note that the Multiple Input One-way Switch M^{k+1} used by Kolaitis and Panttaja [10] is isomorphic to $M^{k,2}$. Hence, $M^{k,n}$ generalizes their construction. The next lemma states the main properties of the switch.

Lemma 9. *For every position $\mathfrak{p}\colon [k] \to [n]$, the following statements hold in the existential (k+1)-pebble game on $M^{k,n}$:*

(i) Spoiler can reach $\{(y^1, y^1_{\mathfrak{p}(1)}), \dots, (y^k, y^k_{\mathfrak{p}(k)})\}$ *from* $\{(x^1, x^1_{\mathfrak{p}(1)}), \dots, (x^k, x^k_{\mathfrak{p}(k)})\}$.

(ii) There are winning strategies $\mathcal{H}^{out}_{(\mathfrak{p},T)}$ *with boundary* $\{(x^i, x^i_0) \mid i \in [k]\}$ *and* $\{(y^i, y^i_{(\mathfrak{p},T)(i)}) \mid i \in [k]\}$ *for all* $T \subseteq [k]$.

(iii) There are winning strategies $\mathcal{H}^{restart}_{(\mathfrak{p},T)}$ *with boundary* $\{(x^i, x^i_{(\mathfrak{p},T)(i)}) \mid i \in [k]\}$ *and* $\{(y^i, y^i_0) \mid i \in [k]\}$ *for all nonempty* $T \subseteq [k]$.

(iv) There exists a critical strategy $\mathcal{H}^{in}_{\mathfrak{p}}$ *with boundary* $\{(x^i, x^i_{\mathfrak{p}(i)}) \mid i \in [k]\} \cup \{(y^i, y^i_0) \mid i \in [k]\}$ *and*

  a) $\mathrm{crit}(\mathcal{H}^{in}_{\mathfrak{p}}) = \bigcup_{t\in[k]} \mathcal{C}^{restart\text{-}crit}_{\mathfrak{p},t} \cup \mathcal{C}^{out\text{-}crit}_{\mathfrak{p}}$,

  b) $\mathcal{C}^{restart\text{-}crit}_{\mathfrak{p},t} \subseteq \mathcal{H}^{restart}_{(\mathfrak{p},\{t\})}$ *and*

  c) $\mathcal{C}^{out\text{-}crit}_{\mathfrak{p}} \subseteq \mathcal{H}^{out}_{\mathfrak{p}}$.

Proof. Let $\mathfrak{p}\colon [k] \to [n]$ be an arbitrary position. We first construct the strategy for Spoiler to prove (i). Starting from position $\{(x^1, x^1_{\mathfrak{p}(1)}), \dots, (x^k, x^k_{\mathfrak{p}(k)})\}$, Spoiler places the $(k{+}1)$st pebble on a^1. Duplicator has to answer with $a^1_{\mathfrak{p}(1),l_1}$ for some $l_1 \in [k]$. Next, Spoiler picks up the pebble from x^1 and puts it on a^2. Duplicator has to answer with $a^2_{\mathfrak{p}(2),l_2}$ for some $l_2 \in [k] \setminus \{l_1\}$. Following that scheme, Spoiler can reach the position $\{(a^1, a^1_{\mathfrak{p}(1),l_1}), \dots, (a^k, a^k_{\mathfrak{p}(k),l_k})\}$ for some $l_1 \neq l_2 \neq \dots \neq l_k$. Now, Spoiler pebbles b^1 with the free pebble and Duplicator has to answer with $b^1_{\mathfrak{p}(1)}$, since every vertex of the form $b^1_{0,l}$ does not form a $(k{+}1)$-clique together with the vertices $a^1_{\mathfrak{p}(1),l_1}, \dots, a^k_{\mathfrak{p}(k),l_k}$. Spoiler picks up the pebble from $a^1_{\mathfrak{p}(1),l_1}$ and puts it on b^2. Duplicator has to answer with $b^2_{\mathfrak{p}(2)}$, because every vertex of the form $b^2_{0,l}$ is not connected to $b^1_{\mathfrak{p}(1)}$. Thus, Spoiler can reach $\{(b^1, b^1_{\mathfrak{p}(1)}), \dots, (b^k, b^k_{\mathfrak{p}(k)})\}$ and from there he reaches $\{(y^1, y^1_{\mathfrak{p}(1)}), \dots, (y^k, y^k_{\mathfrak{p}(k)})\}$ with the same technique.

To prove (ii) let $H^{\mathrm{out}}_{\mathfrak{p}}$ denote the set of total homomorphisms $h^{\mathrm{out}}_{\mathfrak{p},\sigma} = \{(x^i, x^i_0), (a^i, a^i_{\mathfrak{p}(i),\sigma(i)}), (b^i, b^i_{\mathfrak{p}(i)}), (y^i, y^i_{\mathfrak{p}(i)}) \mid i \in [k]\}$, where $\sigma \in S_k$ is some permutation on $[k]$. Furthermore, let $h^{\mathrm{out}}_{(\mathfrak{p},T)} = \{(x^i, x^i_0), (a^i, a^i_0), (b^i, b^i_{0,i}), (y^i, y^i_{(\mathfrak{p},T)(i)}) \mid i \in [k]\}$. Since $h^{\mathrm{out}}_{(\mathfrak{p},T)}$ and all $h^{\mathrm{out}}_{\mathfrak{p},\sigma} \in H^{\mathrm{out}}_{\mathfrak{p}}$ are total,

$$\mathcal{H}^{\mathrm{out}}_{(\mathfrak{p},T)} := \begin{cases} \mathrm{cl}(h^{\mathrm{out}}_{(\mathfrak{p},T)}),\ T \neq \emptyset, \\ \mathrm{cl}(\{h^{\mathrm{out}}_{(\mathfrak{p},\emptyset)}\} \cup H^{\mathrm{out}}_{\mathfrak{p}}),\ \text{else}, \end{cases}$$

are winning strategies for Duplicator satisfying (ii). For all nonempty $T \subseteq [k]$, let $\mathcal{H}^{\mathrm{restart}}_{(\mathfrak{p},T)} := \mathrm{cl}(H^{\mathrm{restart}}_{(\mathfrak{p},T)})$, where $H^{\mathrm{restart}}_{(\mathfrak{p},T)}$ is the set of total homomorphisms h satisfying the constraints $h(x^i) = x^i_{(\mathfrak{p},T)(i)}$ and $h(y^i) = y^i_0$. This set clearly satisfies (iii). As an example let $g \in H^{\mathrm{restart}}_{(\mathfrak{p},\{t\})}$ be the following homomorphism:

$$\begin{aligned} g(x^i) &= x^i_{(\mathfrak{p},\{t\})(i)}, & g(b^i) &= b^i_{0,t}, \\ g(a^t) &= a^t_0, & g(y^i) &= y^i_0, \\ g(a^i) &= a^i_{\mathfrak{p}(i),i},\ i \neq t. \end{aligned}$$

For all $l \in [k]$ and permutations σ on $[k]$ we define partial homomorphisms:

$$\begin{aligned} h^{\mathrm{in}}_{\mathfrak{p},\sigma}(x^i) &= x^i_{\mathfrak{p}(i)} & h^{\mathrm{in}}_{\mathfrak{p},\sigma,l}(x^i) &= x^i_{\mathfrak{p}(i)} \\ h^{\mathrm{in}}_{\mathfrak{p},\sigma}(a^i) &= a^i_{\mathfrak{p}(i),\sigma(i)} & h^{\mathrm{in}}_{\mathfrak{p},\sigma,l}(a^i) &= a^i_{\mathfrak{p}(i),\sigma(i)},\ i \neq \sigma^{-1}(l) \\ h^{\mathrm{in}}_{\mathfrak{p},\sigma}(b^i) &= \text{undefined} & h^{\mathrm{in}}_{\mathfrak{p},\sigma,l}(b^i) &= b^i_{0,l} \\ h^{\mathrm{in}}_{\mathfrak{p},\sigma}(y^i) &= y^i_0 & h^{\mathrm{in}}_{\mathfrak{p},\sigma,l}(y^i) &= y^i_0 \end{aligned}$$

It is easy to check, that $h^{\text{in}}_{\mathfrak{p},\sigma}$ defines a homomorphism from $M^{k,n}_S \setminus B$ to $M^{k,n}_D$, and $h^{\text{in}}_{\mathfrak{p},\sigma,l}$ defines a homomorphism from $M^{k,n}_S \setminus \{a^{\sigma^{-1}(l)}\}$ to $M^{k,n}_D$. For all $\sigma \in S_k$

$$h^{\text{out-crit}}_{\mathfrak{p},\sigma} := \{(a^i, a^i_{\mathfrak{p}(i),\sigma(i)}) \mid i \in [k]\}$$

and for all $\sigma \in S_k$ and $j, t \in [k]$

$$h^{\text{restart-crit}}_{\mathfrak{p},\sigma,j,t} := \{(a^i, a^i_{\mathfrak{p}(i),\sigma(i)}) \mid i \in [k] \setminus \{t\}\} \cup \{(b^j, b^j_{0,\sigma(t)})\}.$$

Now we can define the sets used in (iv):

$$\begin{aligned}
\mathcal{H}^{\text{in}}_{\mathfrak{p}} &= \text{cl}(\{h^{\text{in}}_{\mathfrak{p},\sigma} \mid \sigma \in S_k\} \cup \{h^{\text{in}}_{\mathfrak{p},\sigma,l} \mid \sigma \in S_k, l \in [k]\}),\\
\mathcal{C}^{\text{out-crit}}_{\mathfrak{p}} &= \{h^{\text{out-crit}}_{\mathfrak{p},\sigma} \mid \sigma \in S_k\},\\
\mathcal{C}^{\text{restart-crit}}_{\mathfrak{p},t} &= \{h^{\text{restart-crit}}_{\mathfrak{p},\sigma,j,t} \mid \sigma \in S_k, j \in [k]\},\\
\text{crit}(\mathcal{H}^{\text{in}}_{\mathfrak{p}}) &= \bigcup_{t\in[k]} \mathcal{C}^{\text{restart-crit}}_{\mathfrak{p},t} \cup \mathcal{C}^{\text{out-crit}}_{\mathfrak{p}}.
\end{aligned}$$

First note that $h^{\text{out-crit}}_{\mathfrak{p},\sigma} \subset h^{\text{in}}_{\mathfrak{p},\sigma}$ and $h^{\text{restart-crit}}_{\mathfrak{p},\sigma,j,t} \subset h^{\text{in}}_{\mathfrak{p},\sigma,\sigma(t)}$. Thus, $\text{crit}(\mathcal{H}^{\text{in}}_{\mathfrak{p}}) \subseteq \mathcal{H}^{\text{in}}_{\mathfrak{p}}$. It easily follows from the definitions, that $h^{\text{out-crit}}_{\mathfrak{p},\sigma} \subset h^{\text{out}}_{\mathfrak{p},\sigma}$. Furthermore, every $h^{\text{restart-crit}}_{\mathfrak{p},\sigma,l,t}$ can be extended to a homomorphism in $\mathcal{H}^{\text{restart}}_{(\mathfrak{p},\{t\})}$ by defining the boundary as required and mapping a^t to a^t_0 and b^j to $b^j_{0,\sigma(t)}$ for all $j \in [k]$. This proves statement b) and c) from (iv). It remains to show that $\mathcal{H}^{\text{in}}_{\mathfrak{p}}$ is a critical strategy with critical positions $\text{crit}(\mathcal{H}^{\text{in}}_{\mathfrak{p}})$.

Claim 10. For all $g \in \mathcal{H}^{\text{in}}_{\mathfrak{p}}$ with $|\text{Dom}(g)| \leq k$, either $g \in \text{crit}(\mathcal{H}^{\text{in}}_{\mathfrak{p}})$ or for all $z \in V(M^{k,n}_S)$ there exist an $h \in \mathcal{H}^{\text{in}}_{\mathfrak{p}}$, such that $g \subseteq h$ and $z \in \text{Dom}(h)$.

Proof. As g is a partial homomorphism from $\mathcal{H}^{\text{in}}_{\mathfrak{p}}$, we can fix some $\sigma \in S_k$ and $l \in [k]$ such that $g \subset \{(x^i, x^i_{\mathfrak{p}(i)}), (a^i, a^i_{\mathfrak{p}(i),\sigma(i)}), (b^i, b^i_{0,l}), (y^i, y^i_0) \mid i \in [k]\}$.
Case 1: $|\text{Dom}(g) \cap A| = k$**.** In this case, $g = h^{\text{out-crit}}_{\mathfrak{p},\sigma}$ and hence, $g \in \text{crit}(\mathcal{H}^{\text{in}}_{\mathfrak{p}})$.
Case 2: $|\text{Dom}(g) \cap A| = k-1$**.** If $\text{Dom}(g) \cap B \neq \emptyset$, then $g = h^{\text{restart-crit}}_{\mathfrak{p},\sigma,j,\sigma^{-1}(l)}$ for some $j \in [k]$. Thus, we can assume that $\text{Dom}(g) \cap B = \emptyset$ and show for all z that g satisfies the extension property. If z is the unique element in $A \setminus \text{Dom}(g)$, then $h^{\text{in}}_{\mathfrak{p},\sigma}$ extends g. If $z \in X \cup B \cup Y$, then $h^{\text{in}}_{\mathfrak{p},\sigma,l}$ extends g.
Case 3: $|\text{Dom}(g) \cap A| \leq k-2$**.** Let j_1 and j_2 be two distinct indices such that a^{j_1}, $a^{j_2} \notin \text{Dom}(g)$. Furthermore, we can assume that $\sigma(j_1) = l$. For $z \neq a^{j_1}$ the homomorphism $h^{\text{in}}_{\mathfrak{p},\sigma,l}$ extends g. If $z = a^{j_1}$, then $h^{\text{in}}_{\mathfrak{p},\sigma',l}$ extends g, where $\sigma' := \{(i, \sigma(i)) \mid i \in [k] \setminus \{j_1, j_2\}\} \cup \{(j_1, \sigma(j_2)), (j_2, \sigma(j_1))\}$. ⊣

□

3.3 The Initialization Gadget

The Initialization Gadget $\text{INIT}^{\mathfrak{s}}$ is built of two Multiple Input One-way Switches M^1 and M^2, vertices y in Spoiler's graph and y_1, y_2 in Duplicator's graph (that are colored c_y), and the boundary vertices $x^1, \ldots, x^k$ and $x^1_0, \ldots, x^k_n$ (where all x^i and x^i_l are colored c_{x^i}). The x- and y-vertices are connected to M^1 and M^2 as shown in Figure 5. Lemma 11 provides the strategies on $\text{INIT}^{\mathfrak{s}}$.

Lemma 11. *For every start position $\mathfrak{s}\colon [k] \to [n]$ the following holds in the existential $(k+1)$-pebble game on* $\text{INIT}^{\mathfrak{s}}$*:*

(i) Spoiler can reach $\{(x^i, x^i_{\mathfrak{s}(i)}) \mid i \in [k]\}$ from $\emptyset$.

(ii) There is a winning strategy $\mathcal{I}^{init}$ for Duplicator with boundary $\{(x^i, x^i_{\mathfrak{s}(i)}) \mid i \in [k]\}$.

(iii) For every $\mathfrak{p}\colon [k] \to [n]$ and every $T \subseteq [k]$ there is a critical strategy $\mathcal{I}^{init}_{(\mathfrak{p},T)}$ with boundary $\{(x^i, x^i_{(\mathfrak{p},T)(i)}) \mid i \in [k]\}$ and $\operatorname{crit}(\mathcal{I}^{init}_{(\mathfrak{p},T)}) \subseteq \mathcal{I}^{init}$.

Spoiler's strategy is quite simple. First he pebbles y. Duplicator has to answer with either y_1 or y_2. Then Spoiler can reach $\{(x^i, x^i_{\mathfrak{s}(i)}) \mid i \in [k]\}$ by pebbling through either M^1 or M^2. To construct the strategies for Duplicator, we can combine the strategies of the switches M^1 and M^2 such that she plays an input-strategy on one switch and a restart- or output-strategy on the other switch. Assume that Spoiler reaches a critical position on the switch where Duplicator plays the input-strategy, say M^1. Duplicator can now flip the strategies such that she plays a restart- or output-strategy on M^1, depending on which kind of critical position Spoiler has reached, and an input-strategy on M^2. A detailed proof of Lemma 11 is given in the appendix.

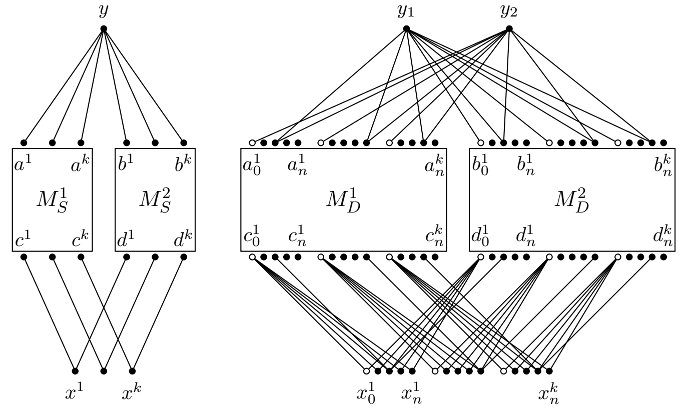

Figure 5: The initialization gadget INIT$^{\mathfrak{s}}$ (for $k = 3$, $n = 4$, $\mathfrak{s}(1) = 2$, $\mathfrak{s}(2) = 4$, $\mathfrak{s}(3) = 3$.)

3.4 The Choice Gadget

The boundary of the Choice Gadget consists of input vertices $x^1, \ldots, x^k$ in Spoiler's graph and $x^1_0, \ldots, x^k_n$ in Duplicator's graph. These vertices will be identified with y-vertices in the final graph. The output vertices are of the form $(y_q)^1, \ldots, (y_q)^k$ and $(y_q)^1_0, \ldots, (y_q)^k_n$ for all $q \in [m]$ and will be connected to the Rule Gadgets $RD(r_q)$. This gadget enables Spoiler to reach positions $((y_q)^i, (y_q)^i_{\mathfrak{p}(i)})$ from $(x^i, x^i_{\mathfrak{p}(i)})$ but Duplicator can choose the desired $q \in [m]$. This choice will later coincide with the rule $r_q \in R$ Player 2 chooses in the KAI-game when position $\mathfrak{p}$ is pebbled. The *Choice Gadget* C^m is defined as follows:

$$
\begin{aligned}
V(C^m_S) &= \{x^i, a^i \mid i \in [k]\} \cup \{(y_q)^i \mid i \in [k], q \in [m]\},\\
E(C^m_S) &= \big\{\{x^i, a^i\} \mid i \in [k]\big\} \cup\\
&\qquad \big\{\{a^i, (y_q)^i\} \mid i \in [k], q \in [m]\big\}\\
&\qquad \cup \big\{\{a^i, a^j\} \mid i, j \in [k];\ i \neq j\big\},\\
V(C^m_D) &= \{x^i_l \mid i \in [k], 0 \leq l \leq n\} \cup \{a^i_0 \mid i \in [k]\}\\
&\qquad \{a^i_{l,q} \mid i \in [k], l \in [n], q \in [m]\}\\
&\qquad \cup \{(y_q)^i_l \mid i \in [k], q \in [m], 0 \leq l \leq n\},\\
E(C^m_D) &= \big\{\{x^i_0, a^i_0\} \mid i \in [k]\big\}\\
&\qquad \cup \big\{\{x^i_l, a^i_{l,q}\} \mid i \in [k], l \in [n], q \in [m]\big\}\\
&\qquad \cup \big\{\{a^i_0, (y_q)^i_0\} \mid i \in [k], q \in [m]\big\}
\end{aligned}
$$

$$
\begin{aligned}
&\cup \big\{\{a^i_{l,q}, (y_q)^i_l\} \mid i \in [k]; l \in [n]; q \in [m]\big\}\\
&\cup \big\{\{a^i_{l,q}, (y_{q'})^i_0\} \mid i \in [k]; l \in [n]; q, q' \in [m]; q \neq q'\big\}\\
&\cup \big\{\{a^i_{l,q}, a^j_{l',q}\} \mid i, j \in [k]; i \neq j; l, l' \in [n]; q \in [m]\big\}.
\end{aligned}
$$

Furthermore, for all $i \in [k]$, all vertices a^i and $a^i_{l,q}$ are colored with the unique color c_{a^i}, and for all $i \in [k]$, $q \in [m]$ the vertices $(y_q)^i$ and $(y_q)^i_l$ are colored with the unique color $c_{(y_q)^i}$. One partition of C^m is shown in Figure 6.

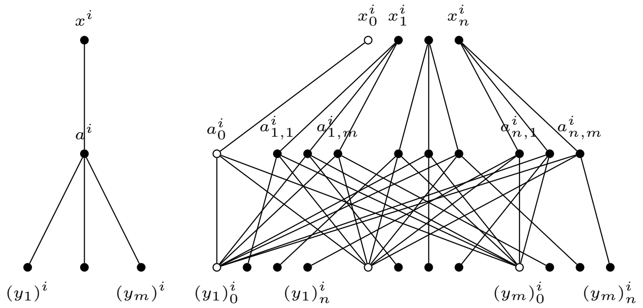

Figure 6: The i-th partition of the gadget C^m.

Lemma 12. *In the existential (k+1)-pebble game on C^m,*

(i) for every $\mathfrak{p}\colon [k] \to [n]$ Spoiler can reach the position $\{((y_l)^i, (y_l)^i_{\mathfrak{p}(i)}) \mid i \in [k]\}$ from $\{(x^i, x^i_{\mathfrak{p}(i)}) \mid i \in [k]\}$ for one $l \in [m]$, and

(ii) for every $l \in [m]$, $\mathfrak{p}\colon [k] \to [n]$, $T \subseteq [k]$, Duplicator has a winning strategy $\mathcal{C}^l_{(\mathfrak{p},T)}$ with boundary $\{(x^i, x^i_{(\mathfrak{p},T)(i)}) \mid i \in [k]\} \cup \{((y_q)^i, (y_q)^i_0) \mid i \in [k], q \in [m] \setminus \{l\}\} \cup \{((y_l)^i, (y_l)^i_{(\mathfrak{p},T)(i)}) \mid i \in [k]\}$.

Proof. We first present Spoiler's strategy (i). Starting from position $\{(x^i, x^i_{\mathfrak{p}(i)}) \mid i \in [k]\}$ Spoiler puts the $(k+1)$st pebble on a^1 and Duplicator has to response with $a^1_{\mathfrak{p}(1),l}$ for one $l \in [m]$ he can choose. Then Spoiler picks up the pebble from x^1 and puts it on a^2. Duplicator has to response with $a^2_{\mathfrak{p}(2),l}$, since this vertex is the only one adjacent to $x^2_{\mathfrak{p}(2)}$ and $a^1_{\mathfrak{p}(1),l}$. Next, Spoiler picks up the pebble from x^2 and puts it on a^3 and so on. Thus, Spoiler reaches $\{(a^i, a^i_{\mathfrak{p}(i),l}) \mid i \in [k]\}$. From there he puts the $(k+1)$st pebble on $(y_l)^1$, and Duplicator has to answer with $(y_l)^1_{\mathfrak{p}(1)}$, since this vertex is the only one colored in the same color as $(y_l)^1$ and adjacent to $a^1_{\mathfrak{p}(1),l}$. Now Spoiler picks up the pebble from a^1 and puts it on $(y_l)^2$, Duplicator has to answer with $(y_l)^2_{\mathfrak{p}(2)}$. Following that strategy Spoiler can reach $\{((y_l)^i, (y_l)^i_{\mathfrak{p}(i)}) \mid i \in [k]\}$. Duplicator's strategy (ii) is simply defined as $\mathcal{C}^l_{(\mathfrak{p},T)} = \mathrm{cl}(h^l_{(\mathfrak{p},T)})$, where $h^l_{(\mathfrak{p},T)}$ is the following total homomorphism on C^m:

$$
h^l_{(\mathfrak{p},T)}(z) := \begin{cases}
x^i_{(\mathfrak{p},T)(i)}, & \text{if } z = x^i,\\
a^i_{\mathfrak{p}(i),l}, & \text{if } z = a^i, i \notin T,\\
a^i_0, & \text{if } z = a^i, i \in T,\\
(y_l)^i_{(\mathfrak{p},T)(i)}, & \text{if } z = (y_l)^i,\\
(y_q)^i_0, & \text{if } z = (y_q)^i,\ q \in [m] \setminus \{l\}.
\end{cases}
$$

□

3.5 Proof of the Main Lemma

Proof of Lemma 4. It remains to construct winning strategies for Spoiler and Duplicator on the colored graphs G_S and G_D. First, we develop the winning strategy for Spoiler. Hence assume that Player 1 has a winning strategy in the k-pebble KAI-game. By Lemma 11 Spoiler can reach position $\{(\mathsf{x}^i, \mathsf{x}^i_{\mathfrak{s}(i)}) \mid i \in [k]\}$. Let r be the first rule Player 1 chooses in the KAI-game. By Lemma 7 Spoiler can reach $\{(y^i, y^i_{r(\mathfrak{s})(i)}) \mid i \in [k]\}$, where the y vertices are the output-boundary of the rule gadget $RS(r)$. By Lemma 9 he can pebble through the switch and reach $\{(\mathsf{y}^i, \mathsf{y}^i_{r(\mathfrak{s})(i)}) \mid i \in [k]\}$. Let $\mathfrak{p} := r(\mathfrak{s})$. If $\mathfrak{p}$ maps some pebble to γ, then Spoiler wins the game. Otherwise (by Lemma 12) Spoiler can reach position $\{((y_q)^i, (y_q)^i_{\mathfrak{p}(i)}) \mid i \in [k]\}$ for one $q \in [m]$ of Duplicator's choice at the output of the choice gadget and the input of $\mathrm{RD}(r_q)$. If $r_q \notin \mathrm{appl}(\mathfrak{p})$, then Spoiler wins immediately by Lemma 8 (especially Spoiler wins if $\mathrm{appl}(\mathfrak{p}) = \emptyset$ and Player 2 cannot make a move). If $r_q \in \mathrm{appl}(\mathfrak{p})$, then Spoiler can reach $\{(\mathsf{x}^i, \mathsf{x}^i_{r_q(\mathfrak{p})(i)}) \mid i \in [k]\}$. Spoiler chooses the next rule according to Player 1's winning strategy and so on. Since Player 1 eventually puts a pebble c on node γ, Spoiler can reach position $(\mathsf{y}^c, \mathsf{y}^c_\gamma)$ and thus he wins the game.

Assume that $\mathcal{K}$ is a winning strategy for Player 2 in the k-pebble KAI-game and let $\mathcal{K}_i := \{\mathfrak{p} \mid (\mathfrak{p}, i) \in \mathcal{K})\}$. For every position $\mathfrak{p} \in \mathcal{K}_2$ let $\mathrm{choice}(\mathfrak{p}) \in [m]$ be the index of the rule r Player 2 chooses according to her winning strategy. Furthermore, let $\mathcal{C}_{\mathbf{0}} := \mathcal{C}^1_{\mathbf{0}} = \cdots = \mathcal{C}^m_{\mathbf{0}}$. Now we construct critical strategies for the whole graph out of smaller critical strategies $\mathcal{F}$ defined on gadgets Q (denoted $\mathcal{F}\langle Q\rangle$) using the $\uplus$-operator and Lemma 6:

$$\begin{aligned}\mathcal{S}^{\mathrm{init}} = \mathcal{I}^{\mathrm{init}} \uplus \mathcal{C}_{\mathbf{0}} \uplus & \\ & \biguplus_{l \in \mathrm{appl}(\mathfrak{s})} \Big(\mathcal{R}_{\mathfrak{s}}\langle RS(r_l)\rangle \uplus \mathcal{H}^{\mathrm{in}}_{r_l(\mathfrak{s})}\langle MS(r_l)\rangle\Big) \uplus \\ & \biguplus_{l \in [m] \setminus \mathrm{appl}(\mathfrak{s})} \Big(\mathcal{R}_{\mathfrak{s}}\langle RS(r_l)\rangle \uplus \mathcal{H}^{\mathrm{restart}}_{(r_l(\mathfrak{s}), T_{r_l})}\langle MS(r_l)\rangle\Big) \uplus \\ & \biguplus_{l \in [m]} \big(\mathcal{R}_{\mathbf{0}}\langle RD(r_l)\rangle \uplus \mathcal{H}^{\mathrm{out}}_{\mathfrak{s}}\langle MD(r_l)\rangle\big)\,.\end{aligned}$$

We define the global critical strategies $\mathcal{S}_{\mathfrak{p}}$ and $\mathcal{S}^{\mathrm{restart}}_{(\mathfrak{p},T)}$ for all $\mathfrak{p} \in \mathcal{K}_1$ and $T \neq \emptyset$:

$$\begin{aligned}\mathcal{S}_{\mathfrak{p}} &= \mathcal{I}^{\mathrm{init}}_{\mathfrak{p}} \uplus \mathcal{C}_{\mathbf{0}} \uplus \\ & \biguplus_{l \in \mathrm{appl}(\mathfrak{p})} \Big(\mathcal{R}_{\mathfrak{p}}\langle RS(r_l)\rangle \uplus \mathcal{H}^{\mathrm{in}}_{r_l(\mathfrak{p})}\langle MS(r_l)\rangle\Big) \uplus \\ & \biguplus_{l \in [m] \setminus \mathrm{appl}(\mathfrak{p})} \Big(\mathcal{R}_{\mathfrak{p}}\langle RS(r_l)\rangle \uplus \mathcal{H}^{\mathrm{restart}}_{(r_l(\mathfrak{p}), T_{r_l})}\langle MS(r_l)\rangle\Big) \uplus \\ & \biguplus_{l \in [m]} \big(\mathcal{R}_{\mathbf{0}}\langle RD(r_l)\rangle \uplus \mathcal{H}^{\mathrm{out}}_{\mathfrak{p}}\langle MD(r_l)\rangle\big)\,, \\ \mathcal{S}^{\mathrm{restart}}_{(\mathfrak{p},T)} &= \mathcal{I}^{\mathrm{init}}_{(\mathfrak{p},T)} \uplus \mathcal{C}_{\mathbf{0}} \uplus \\ & \biguplus_{l \in [m]} \big(\mathcal{R}_{(\mathfrak{p},T)}\langle RS(r_l)\rangle \uplus \mathcal{H}^{\mathrm{restart}}_{(r_l(\mathfrak{p}), T \cup T_{r_l})}\langle MS(r_l)\rangle \uplus \\ & \quad \mathcal{R}_0\langle RD(r_l)\rangle \uplus \mathcal{H}^{\mathrm{out}}_{(\mathfrak{p},T)}\langle MD(r_l)\rangle\big).\end{aligned}$$

Furthermore, for all $\mathfrak{p} \in \mathcal{K}_2$ and $T \neq \emptyset$ let $\mathcal{D}_{\mathfrak{p}}$ and $\mathcal{D}^{\mathrm{restart}}_{(\mathfrak{p},T)}$ be the following global critical strategies:

$$\mathcal{D}_{\mathfrak{p}} = \mathcal{I}^{\mathrm{init}}_{\mathbf{0}} \uplus \mathcal{C}^{\mathrm{choice}(\mathfrak{p})}_{\mathfrak{p}} \uplus$$

$$
\begin{aligned}
&\biguplus_{l\in[m]} \big(\mathcal{R}_{\mathbf{0}}\langle RS(r_l)\rangle \uplus \mathcal{H}^{\text{out}}_{\mathfrak{p}}\langle MS(r_l)\rangle\big)\uplus\\
&\qquad\biguplus_{l\in[m]\setminus\{\text{choice}(\mathfrak{p})\}} \big(\mathcal{R}_{\mathbf{0}}\langle RD(r_l)\rangle \uplus \mathcal{H}^{\text{restart}}_{\mathbf{0}}\langle MD(r_l)\rangle\big)\uplus\\
&\mathcal{R}_{\mathfrak{p}}\langle RD(r_{\text{choice}(\mathfrak{p})})\rangle \uplus \mathcal{H}^{\text{in}}_{r_{\text{choice}(\mathfrak{p})}(\mathfrak{p})}\langle MD(r_{\text{choice}(\mathfrak{p})})\rangle,
\end{aligned}
$$

$$
\begin{aligned}
\mathcal{D}^{\text{restart}}_{(\mathfrak{p},T)} = {}&\mathcal{I}^{\text{init}}_{\mathbf{0}} \uplus \mathcal{C}^{\text{choice}(\mathfrak{p})}_{(\mathfrak{p},T)}\uplus\\
&\biguplus_{l\in[m]} \big(\mathcal{R}_{\mathbf{0}}\langle RS(r_l)\rangle \uplus \mathcal{H}^{\text{out}}_{(\mathfrak{p},T)}\langle MS(r_l)\rangle\big)\uplus\\
&\qquad\biguplus_{l\in[m]\setminus\{\text{choice}(\mathfrak{p})\}} \big(\mathcal{R}_{\mathbf{0}}\langle RD(r_l)\rangle \uplus \mathcal{H}^{\text{restart}}_{\mathbf{0}}\langle MD(r_l)\rangle\big)\uplus\\
&\mathcal{R}_{(\mathfrak{p},T)}\langle RD(r_{\text{choice}(\mathfrak{p})})\rangle\uplus\\
&\quad\mathcal{H}^{\text{restart}}_{(r_{\text{choice}(\mathfrak{p})}(\mathfrak{p}),T)}\langle MD(r_{\text{choice}(\mathfrak{p})})\rangle.
\end{aligned}
$$

Before we formally state Duplicator's winning strategy, we briefly describe these critical strategies. First, the only critical positions of $\mathcal{S}^{\text{restart}}_{(\mathfrak{p},T)}$ and $\mathcal{D}^{\text{restart}}_{(\mathfrak{p},T)}$ are inside the Initialization Gadget and contained in $\mathcal{S}^{\text{init}}$. Thus, this is a good situation for Duplicator, since Spoiler has to restart the game by playing on the Initialization Gadget. At the beginning of the game Duplicator plays according to the strategy $\mathcal{S}^{\text{init}}$, where position $\mathfrak{s}$ is on the x-vertices. The only critical positions of that strategy are inside of the $MS(r)$-gadgets for rules r applicable to $\mathfrak{s}$. If Spoiler pebbles some restart-critical positions there, then Duplicator can switch to $\mathcal{S}^{\text{restart}}_{(\mathfrak{s},T)}$. If Spoiler pebbles an output-critical position on $MS(r)$, then Duplicator can switch to strategy $\mathcal{D}_{r(\mathfrak{s})}$ where position $\mathfrak{p} = r(\mathfrak{s})$ is on the y-vertices. The only critical positions now are inside of the switch $MD(r_{\text{choice}(\mathfrak{p})})$ and inside the Initialization Gadget. If Spoiler pebbles a restart-critical position on $MD(r_{\text{choice}(\mathfrak{p})})$, then Duplicator sticks to $\mathcal{D}^{\text{restart}}_{(\mathfrak{p},T)}$. If Spoiler pebbles an output-critical position, then Duplicator chooses strategy $\mathcal{S}_{\mathfrak{p}'}$, where $\mathfrak{p}' = r_{\text{choice}(\mathfrak{p})}(\mathfrak{p})$. The critical positions from $\mathcal{S}_{\mathfrak{p}'}$ are within the Initialization Gadget and the switches $MS(r)$ for rules r applicable to $\mathfrak{p}'$. Thus, combining all these critical strategies allows Duplicator to play forever. Now we define the winning strategy for Duplicator:

$$
\begin{aligned}
\mathcal{H} = \mathcal{S}^{\text{init}} \cup &\bigcup_{\mathfrak{p}\in\mathcal{K}_1, T\subseteq[k], T\neq\emptyset} \left(\mathcal{S}_{\mathfrak{p}} \cup \mathcal{S}^{\text{restart}}_{(\mathfrak{p},T)}\right)\\
\cup &\bigcup_{\mathfrak{p}\in\mathcal{K}_2, T\subseteq[k], T\neq\emptyset} \left(\mathcal{D}_{\mathfrak{p}} \cup \mathcal{D}^{\text{restart}}_{(\mathfrak{p},T)}\right).
\end{aligned}
$$

Since $\mathcal{H}$ is a union of critical strategies, it suffices by Lemma 5 to show that for each critical strategy $\mathcal{G}$ and each partial homomorphism $h \in \text{crit}(\mathcal{G})$ there is a critical strategy $\mathcal{F}$ such that $h \in \hat{\mathcal{F}}$. From the definition of the global critical strategies and the properties of the partial critical strategies they contain, it follows

$$
\begin{aligned}
\text{crit}(\mathcal{S}^{\text{restart}}_{(\mathfrak{p},T)}) &\subseteq \hat{\mathcal{S}}^{\text{init}},\\
\text{crit}(\mathcal{D}^{\text{restart}}_{(\mathfrak{p},T)}) &\subseteq \hat{\mathcal{S}}^{\text{init}},\\
\text{crit}(\mathcal{S}_{\mathfrak{p}}) &\subseteq \hat{\mathcal{S}}^{\text{init}} \cup \bigcup_{T\neq\emptyset} \hat{\mathcal{S}}^{\text{restart}}_{(\mathfrak{p},T)} \cup \bigcup_{l\in\text{appl}(\mathfrak{p})} \hat{\mathcal{D}}_{r_l(\mathfrak{p})},\\
\text{crit}(\mathcal{D}_{\mathfrak{p}}) &\subseteq \hat{\mathcal{S}}^{\text{init}} \cup \hat{\mathcal{S}}_{r_{\text{choice}(\mathfrak{p})}(\mathfrak{p})} \cup \bigcup_{T\neq\emptyset} \hat{\mathcal{D}}^{\text{restart}}_{(\mathfrak{p},T)},\\
\text{crit}(\mathcal{S}^{\text{init}}) &\subseteq \bigcup_{l\in\text{appl}(\mathfrak{s})} \hat{\mathcal{D}}_{r_l(\mathfrak{s})}.
\end{aligned}
$$

From the definition of $\mathcal{H}$ and the properties of the KAI-game winning strategy $\mathcal{K}$, it follows, that if a global critical strategy mentioned in the left hand side of the above inclusions is a strategy in $\mathcal{H}$, then so are all strategies on the right hand side. This concludes the proof of Lemma 4 for colored simple graphs G_S and G_D. □

3.6 Getting Rid of the Colors

As in [10] our construction involves $|V(G_S)|$ unary predicates. To settle the complexity of deciding whether Spoiler has a winning strategy in the existential k-pebble game on σ-structures for fixed finite signatures σ, we use the following construction to switch from colored simple graphs to directed graphs. Let $P_1, \ldots, P_w$ be the colors used in the graphs G_S and G_D, and P an additional color. We introduce $w+1$ vertices $d_0, \ldots, d_w$ that are colored P in both graphs G_S and G_D. For every $1 \leq i < w$ and each vertex x colored P_i, there are directed edges (d_i, x) and (x, d_{i+1}). Furthermore, there are directed edges from vertices colored P_w to d_w and one from d_1 to d_0. Now we can delete the colors $P_1, \ldots, P_w$ without giving Duplicator more freedom. That is, we replace the requirement "$x \in P_i$" by the statement "there exists an alternating directed path (where every second element is colored P) of length $2i-1$ to a vertex colored P and having an out-neighbor colored P". This can easily be checked by Spoiler using two pebbles. To get rid of the remaining color P we add a loop $E(x,x)$ for every vertex $x \in P$.

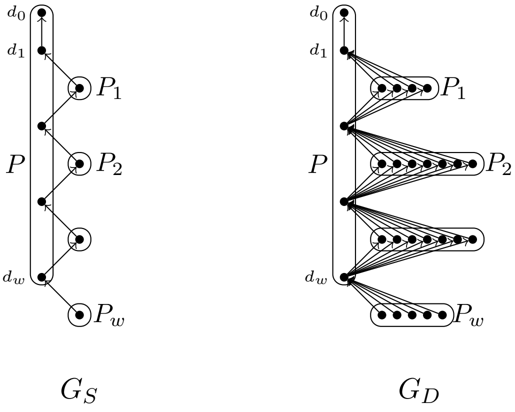

Figure 7: Colors $P_1, \ldots, P_w$ can now be deleted.

4 Conclusion

We proved an $\omega(n^{\frac{(k-3)}{12}})$ lower bound for deciding which player can win the existential k-pebble game on two given finite relational structures and got an $\omega(n^{\frac{(k-3)}{12}})$ lower bound for the k-consistency test as a consequence. Furthermore, the lower bound also applies to the k-pebble game that characterizes the expressive power of the existential k-variable fragment (where negation is allowed in front of atomic formulas).

The parameterized complexity of the whole k-variable first-order logic L^k and the counting logic C^k is wide open. It is even not known if it is W[1]-hard to decide if two given finite relational structures can be distinguished by an L^k (C^k) sentence. Regarding the classical complexity, L^k-equivalence as well as C^k-equivalence is complete for polynomial time [8], but it is an open problem whether EXPTIME-completeness holds when k is part of the input.

References

[1] A. Adachi, S. Iwata, and T. Kasai, “Some combinatorial game problems require $\Omega(n^k)$ time,” *J. ACM*, vol. 31, Mar. 1984.

[2] A. Atserias, A. A. Bulatov, and V. Dalmau, “On the power of k -consistency,” in *Proc. ICALP’07*, 2007, pp. 279–290.

[3] M. C. and Cooper, “An optimal k-consistency algorithm,” *Artificial Intelligence*, vol. 41, no. 1, pp. 89 – 95, 1989.

[4] V. Dalmau, P. G. Kolaitis, and M. Y. Vardi, “Constraint satisfaction, bounded treewidth, and finite-variable logics,” in *Proc. CP’02*, 2002, pp. 310–326.

[5] R. G. Downey and M. R. Fellows, *Parameterized Complexity.* Springer-Verlag, 1999.

[6] T. Feder and M. Y. Vardi, “The computational structure of monotone monadic snp and constraint satisfaction: A study through datalog and group theory,” *SIAM Journal on Computing*, vol. 28, no. 1, pp. 57–104, 1998.

[7] S. Gaspers and S. Szeider, “The parameterized complexity of local consistency,” in *Proc. CP’11*, 2011, pp. 302–316.

[8] M. Grohe, “Equivalence in finite-variable logics is complete for polynomial time,” in *In Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science*, 1996, pp. 264–273.

[9] T. Kasai, A. Adachi, and S. Iwata, “Classes of pebble games and complete problems,” *SIAM J. Comput.*, vol. 8, no. 4, pp. 574–586, 1979.

[10] P. G. Kolaitis and J. Panttaja, “On the complexity of existential pebble games,” in *Proc. CSL’03*, 2003, pp. 314–329.

[11] P. G. Kolaitis and M. Y. Vardi, “On the expressive power of datalog: Tools and a case study,” *J. Comput. Syst. Sci.*, vol. 51, no. 1, pp. 110–134, 1995.

[12] ——, “Conjunctive-query containment and constraint satisfaction,” *J. Comput. Syst. Sci.*, vol. 61, pp. 302–332, Oct. 2000.

[13] ——, “A game-theoretic approach to constraint satisfaction,” in *Proc AAAI/IAAI’00*, 2000, pp. 175–181.

Appendix

Proof of Lemma 11. We first develop the strategy for Spoiler (i). Spoiler first pebbles y. Duplicator has to response either with y_1 or with y_2. Depending on Duplicator’s choice, Spoiler can reach either $\{(a^i, a^i_{\mathfrak{s}(i)}) \mid i \in [k]\}$ or $\{(b^i, b^i_{\mathfrak{s}(i)}) \mid i \in [k]\}$. By Lemma 9.(i) Spoiler reaches $\{(c^i, c^i_{\mathfrak{s}(i)}) \mid i \in [k]\}$ ($\{(d^i, d^i_{\mathfrak{s}(i)}) \mid i \in [k]\}$) and from there he can reach the position $\{(x^i, x^i_{\mathfrak{s}(i)}) \mid i \in [k]\}$.

In order to construct the winning strategy for Duplicator, we let $R, T \subseteq [k]$ and $\mathfrak{p} \colon [k] \to [n]$ and define winning strategies on the boundary of the switches together with the x- and y-vertices. As abbreviation we drop the “for all $i \in [k]$”.

$$\mathcal{G}^1_R := \mathrm{cl}(\{(y, y_1)\} \cup \{(a^i, a^i_{\mathfrak{s}(i)})\} \cup \{(b^i, b^i_{(\mathfrak{s},R)(i)})\})$$
$$\mathcal{G}^2_R := \mathrm{cl}(\{(y, y_2)\} \cup \{(b^i, b^i_{\mathfrak{s}(i)})\} \cup \{(a^i, a^i_{(\mathfrak{s},R)(i)})\})$$

$$
\begin{aligned}
\mathcal{K}_{(\mathfrak{p},T)} &:= \mathrm{cl}(\{(c^i, c^i_0)\} \cup \{(d^i, d^i_0)\} \cup \{(x^i, x^i_{(\mathfrak{p},T)(i)})\}) \\
\mathcal{K}^{\text{out-1}} &:= \mathrm{cl}(\{(c^i, c^i_{\mathfrak{s}(i)})\} \cup \{(d^i, d^i_0)\} \cup \{(x^i, x^i_{\mathfrak{s}(i)})\}) \\
\mathcal{K}^{\text{out-2}} &:= \mathrm{cl}(\{(c^i, c^i_0)\} \cup \{(d^i, d^i_{\mathfrak{s}(i)})\} \cup \{(x^i, x^i_{\mathfrak{s}(i)})\})
\end{aligned}
$$

By Lemma 6 we get the following critical strategies on $\mathrm{INIT}^{\mathfrak{s}}$.

$$
\begin{aligned}
\mathcal{I}^{\text{in-1}}_{R,(\mathfrak{p},T)} &:= \mathcal{G}^1_R \uplus \mathcal{H}^{\text{in}}_{\mathfrak{s}}\langle M^1\rangle \uplus \mathcal{H}^{\text{restart}}_{(\mathfrak{s},R)}\langle M^2\rangle \uplus \mathcal{K}_{(\mathfrak{p},T)} \\
\mathcal{I}^{\text{in-2}}_{R,(\mathfrak{p},T)} &:= \mathcal{G}^2_R \uplus \mathcal{H}^{\text{restart}}_{(\mathfrak{s},R)}\langle M^1\rangle \uplus \mathcal{H}^{\text{in}}_{\mathfrak{s}}\langle M^2\rangle \uplus \mathcal{K}_{(\mathfrak{p},T)} \\
\mathcal{I}^{\text{init-1}} &:= \mathcal{G}^2_{[k]} \uplus \mathcal{H}^{\text{out}}_{\mathfrak{s}}\langle M^1\rangle \uplus \mathcal{H}^{\text{in}}_{\mathfrak{s}}\langle M^2\rangle \uplus \mathcal{K}^{\text{out-1}} \\
\mathcal{I}^{\text{init-2}} &:= \mathcal{G}^1_{[k]} \uplus \mathcal{H}^{\text{in}}_{\mathfrak{s}}\langle M^1\rangle \uplus \mathcal{H}^{\text{out}}_{\mathfrak{s}}\langle M^2\rangle \uplus \mathcal{K}^{\text{out-2}} \\
\mathcal{I}^{\text{init}} &:= \mathcal{I}^{\text{init-1}} \cup \mathcal{I}^{\text{init-2}} \cup \bigcup_{t\in[k]} (\mathcal{I}^{\text{in-1}}_{\{t\},\mathfrak{s}} \cup \mathcal{I}^{\text{in-2}}_{\{t\},\mathfrak{s}}) \\
\mathcal{I}^{\text{init}}_{(\mathfrak{p},T)} &:= \bigcup_{t\in[k]} (\mathcal{I}^{\text{in-1}}_{\{t\},(\mathfrak{p},T)} \cup \mathcal{I}^{\text{in-2}}_{\{t\},(\mathfrak{p},T)})
\end{aligned}
$$

It follows from the definition that $\mathcal{I}^{\text{init}}$ and $\mathcal{I}^{\text{init}}_{(\mathfrak{p},T)}$ have the desired boundary. Furthermore, by Lemma 9.(iv):

$$
\begin{aligned}
\mathrm{crit}(\mathcal{I}^{\text{in-2}}_{R,(\mathfrak{p},T)}) = \mathrm{crit}(\mathcal{I}^{\text{init-1}}) = \mathrm{crit}(\mathcal{H}^{\text{in}}_{\mathfrak{s}}\langle M^2\rangle) \\
&\subseteq \mathcal{H}^{\text{out}}_{\mathfrak{s}}\langle M^2\rangle \cup \bigcup_{t\in[k]} \mathcal{H}^{\text{restart}}_{(\mathfrak{s},\{t\})}\langle M^2\rangle \\
&\subseteq \hat{\mathcal{I}}^{\text{init-2}} \cup \bigcup_{t\in[k]} \hat{I}^{\text{in-1}}_{\{t\},\mathfrak{s}}, \\
\mathrm{crit}(\mathcal{I}^{\text{in-1}}_{R,(\mathfrak{p},T)}) = \mathrm{crit}(\mathcal{I}^{\text{init-2}}) = \mathrm{crit}(\mathcal{H}^{\text{in}}_{\mathfrak{s}}\langle M^1\rangle) \\
&\subseteq \mathcal{H}^{\text{out}}_{\mathfrak{s}}\langle M^1\rangle \cup \bigcup_{t\in[k]} \mathcal{H}^{\text{restart}}_{(\mathfrak{s},\{t\})}\langle M^1\rangle \\
&\subseteq \hat{\mathcal{I}}^{\text{init-1}} \cup \bigcup_{t\in[k]} \hat{I}^{\text{in-2}}_{\{t\},\mathfrak{s}}.
\end{aligned}
$$

Hence, $\mathrm{crit}(\mathcal{I}^{\text{init}}_{(\mathfrak{p},T)}) \subseteq \mathcal{I}^{\text{init}}$ and $\mathcal{I}^{\text{init}}$ is a winning strategy by Lemma 5.

□